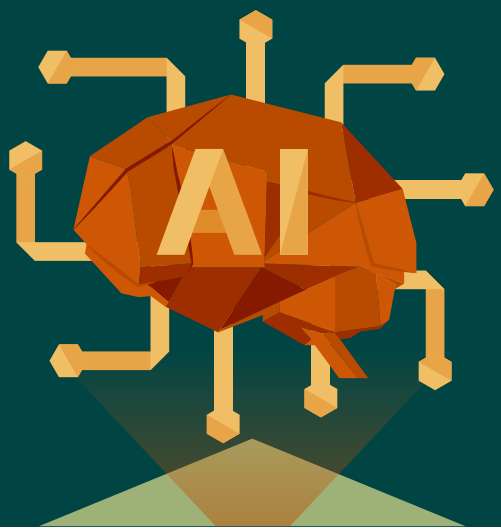
Spectro Cloud
GOVERNMENT

Sovereign AI. On your terms.

Your data, your infrastructure, your choice of models. Spectro Cloud gives state and local government a practical and budget-friendly path to sovereign AI, from private inference for developers to AI factories and secure edge deployments.



AI adoption brings new dependencies

AI can help your teams write software, process documents and serve residents faster. But routing every task to a frontier provider puts critical capabilities outside your control. Pricing changes. Models retire. Access rules shift.

GSA's removal and subsequent restoration of Anthropic in 2026 made that dependency tangible. Even when a restriction does not apply to your agency, it raises a reasonable question: how easily could you change course?

The skills gap is immediate: 76% of states surveyed needed additional staff for AI modeling, according to the 2026 Digital States Survey. Building and maintaining a DIY AI stack adds to an already pressing burden.

Our approach: local control, open choices

We recommend making local AI a capability your agency owns, with the freedom to use hosted services where they add value. With on-prem AI, sensitive workloads can stay wholly within your environment. Your existing frontier services can keep earning their place where policy permits.

Sovereignty means deciding where AI runs, which models it uses and what it can spend. It also means having alternatives when a supplier's terms change.

Spectro Cloud makes that choice practical by packaging the infrastructure and automating its management. Your team can focus on useful applications without needing to become an expert in every layer underneath.

© 2026 Spectro Cloud. All rights reserved.

The right platform for your AI

Local AI, ready to use

PaletteAI Inference Launchpad (PAIL) is a complete software appliance with everything you need to deliver AI locally on your GPU hardware. Developers can use familiar coding assistants and an OpenAI-compatible API with local models. It runs completely airgapped by default, and admins have full, granular control over usage, permissions and spend quotas.

Your AI factory, faster

For larger deployments like AI factories, **PaletteAI VerteX** is the answer. It simplifies repeatable infrastructure deployment with prevalidated stacks. Shared GPU pools and workload-aware allocation improve utilization across teams. Developers get on-demand capacity while IT manages the environment centrally.

AI at the edge

Many AI workloads belong at the edge: on the streets, in vehicles, and across distributed sites. PaletteAI VerteX makes it easy to deploy, secure and operate edge AI infrastructure at scale, with centralized fleet management and local policy enforcement. It also runs containers and, with **Virtual Machine Orchestrator**, VMs: meaning it provides a common operational foundation for new AI services and the applications you already depend on.

Results you can count on

- **Data sovereignty.** Local AI processing keeps sensitive inference within infrastructure you control, with explicit choices about external access.
- **Budget discipline.** Owned capacity, shared GPUs and usage limits make AI spending easier to plan and account for.
- **Faster delivery.** Packaged stacks and repeatable deployments reduce integration work and give scarce technical staff time back.
- **Room to evolve.** Open infrastructure supports changing models, tools and deployment needs without making every new project a platform rebuild.

Built on government experience

Spectro Cloud has spent years supporting government AI infrastructure from data centers to disconnected edge locations, with US-based, US-citizen support available 24×7.

Our VerteX SaaS platform is GovRAMP Authorized at Moderate and FedRAMP Class C (Moderate) certified. Our Cryptographic Library is FIPS 140-3 Level 1 validated (#5061). Our government team helps apply that experience and evidence to your deployment.

GOVRAMP MEMBER

FedRAMP

FIPS 140-3™

Get started today

Connect with us to discuss your AI projects, whether you're looking to reduce dependency on the frontier, or scaling out a major infrastructure build.

Talk to our government team: spectrocloud.com/government

Spectro Cloud GOVERNMENT

© 2026 Spectro Cloud. All rights reserved. Spectro Cloud, Palette, and the respective logos are trademarks or registered trademarks of Spectro Cloud in the US and other countries. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Spectro Cloud and any other company. File: Spectro Cloud sovereign AI_solutions brief_20261102